

TORTURE IN PRINCIPLE AND IN PRACTICE

Jeff McMahan

I. AGAINST MORAL ABSOLUTISM

Those of us who oppose torture, and who are acutely conscious of the grave wrongs being committed in our name by our current government, had better be clear and convincing about the basis of our opposition. While I admire the spirit of Ben Juratowitch's essay, I cannot accept its arguments.¹ I believe that the case against torture cannot plausibly take an absolutist form and that effective opposition to torture is ill-served by appeals to unexplicated and ultimately unserviceable notions such as that torture violates the victim's human dignity and undermines the perpetrator's humanity. We fail to take the problem of torture sufficiently seriously if we treat it as a simple matter of civilization versus barbarism, or a choice between respect for human dignity and a collapse into moral degradation and defilement.

In this section I will explain in a quite general way why I believe that absolute prohibitions of act-types such as torture and killing are unacceptable. In the second section, I will elucidate the grounds on which torture can be morally permissible in principle. In the third and final section, I will argue that the moral justifiability of torture in principle is virtually irrelevant in practice and that it is morally necessary that the law, both domestic and international, should prohibit the practice of torture absolutely—that is, without exceptions.

One surprising feature of the debate about torture is that a great many opponents of torture adopt, or present themselves as adopting, the view that torture is in principle absolutely prohibited by morality.² Nothing, on this view, could ever justify torture. What is surprising about this is that most of these people seem to reject absolutism in all other areas of morality. Most of them, for example, are not absolutists about killing. And it is easy to see why if we survey the more prominent variants of the view that killing is absolutely prohibited by morality.

One view is that it is absolutely impermissible to kill an innocent person. Stated this simply, however, such a view is doubtfully coherent, since it seems possible that there could be cases in which whatever a person does, she will kill an innocent person. So perhaps this first version of absolutism about killing should instead be

that whenever there is an option that does not involve killing an innocent person, it is absolutely prohibited to kill an innocent person.

Note that this view applies to all instances of killing, whether the killing is intended or merely foreseen but unintended. Because of this, it provides the basis for what I think is the most plausible version of pacifism, which claims that war is invariably wrong because it always involves the killing of innocent people. This is not an absolutist form of pacifism because it does not rule out a war that would not involve any killing of innocent people—for example, a war at sea or in outer space, assuming that combatants on neither side were innocent in the relevant sense. But it is a form of pacifism because it rules out all wars that we are ever likely to fight.

The problem with this form of absolutism about killing is that it attributes excessive weight to the significance of the distinction between killing and letting die, and no weight at all to the distinction between intended killing and killing that is unintended though foreseen. Suppose that there is a single military base from which a group of bombers will fly to drop bombs on a city in which 100,000 innocent people live. Suppose further that one can save all these people by destroying the base before the bombers can take off, but that in doing so one will unavoidably kill one innocent person as a side effect. The view that an avoidable killing of an innocent person can never be permissible implies that one ought to allow the 100,000 innocent people to be killed. Although I accept that the distinction between what we do and what we allow to happen has moral significance, it is hard to believe that it is sufficiently significant to make the destruction of the base morally impermissible.

A more plausible absolutist view about killing is that it is absolutely impermissible to kill innocent people *intentionally*. Yet most of us reject this view on intuitive grounds. Suppose that, for whatever reason, the only means of preventing the destruction of the city with its 100,000 innocent inhabitants is to kill one innocent person. To suppose that it would be permissible to kill this person as a side effect, as in the previous version of the example, but absolutely impermissible to kill him intentionally, is to attribute excessive significance to intention. Although I accept that in general it is more seriously objectionable to harm a person intentionally than to cause him the same harm foreseeably but unintentionally, it is hard to believe that what an agent intends in acting can make as much difference as this form of absolutism assumes.

This is, of course, merely an appeal to intuition. But there is a more serious problem for this form of absolutism about killing. (The same problem arises for the previous version as well.) Assume that innocence is all-or-nothing, that is, that innocence is not a matter of degree. And assume further that what it is for a person to be innocent, in the sense relevant to the permissibility of killing, is that the person bears no moral responsibility for a wrong, such as a threat of wrongful harm, that might be prevented or corrected by killing him. (I believe that this is the correct substantive sense of the term in this context, though I cannot argue for that here.) Given these assumptions, noninnocence must be a matter of degree, since moral

responsibility comes in degrees. Next consider two people, each of whom poses a threat to a large number of innocent people. One bears no moral responsibility whatsoever for the threat he poses (he may have been involuntarily administered a drug that has rendered him irresistibly susceptible to suggestion), while the other bears only the slightest possible degree of responsibility for the threat he poses. The view that it is absolutely impermissible intentionally to kill an innocent person, but not necessarily impermissible to kill a relevantly noninnocent person, prohibits the killing of the first threatening person, no matter how much harm he will otherwise cause, but allows that it may be permissible to kill the second, even if the harm he would cause would be of a substantially lesser magnitude.

Some may think that this objection is easy to evade because on what they regard as a more plausible conception of innocence, noninnocence is also all-or-nothing. For example, many people believe that to be innocent in war is simply to pose no threat to others, so that to be noninnocent is to pose a threat; and a person either poses a threat or he does not. But whatever conception of innocence one adopts, there remains a similar problem: the problem of uncertainty. Suppose that one could save many people's lives by killing one person, but that one cannot be certain whether this person is innocent in the relevant sense. On some conceptions of innocence it may be hard to imagine cases in which this is true. But I suspect that such cases are always possible. If, for example, we accept the common view that a person is noninnocent in the relevant sense if he poses a threat to others, we can imagine a case in which we are uncertain whether a person actually poses a threat but are confident that, if he does, killing him will eliminate the threat, and that, if he does not, killing him will nevertheless eliminate the threat in a different way. Suppose that in such a case it is reasonable to believe that there is a 60 percent probability that he does not pose a threat and is therefore innocent. It is hard to see how a theory that implies that it could be permissible to kill him could be said to assert an absolute prohibition of the intentional killing of the innocent. Yet the same seems true even as we progressively lower the probability that he is innocent. Even if there is only a 5 percent probability that he is innocent, how can a theory that implies that it is permissible to kill him count as absolutely prohibiting the intentional killing of the innocent?

It seems that an absolutist prohibition of the intentional killing of the innocent must insist that the intentional killing of a person can be permissible only if it is certain that the person is noninnocent. Yet in practice this would be tantamount to an absolute prohibition of the intentional killing of persons, whether innocent or noninnocent, since one can never in practice be *certain* of a person's noninnocence.

Some pacifists do claim that the intentional killing of any person is absolutely prohibited. So do some others who are not pacifists because they believe that it is possible to participate in war intending only to incapacitate one's enemies, though foreseeing that one's means of incapacitating them may also kill them as a side effect. But the price of accepting this view is the rejection of fundamental principles

of justice. If a man is on the verge of killing an innocent child and the only way one can prevent him from doing it is to kill him, it is permissible as a matter of justice to kill him. By his own voluntary action he has made it the case that either he or the child will be killed. It is a matter of justice in the distribution of harm that he should pay the cost of his own wrongful action. Given what he has done, he cannot reasonably object to being killed, and he will not be wronged if he is killed.

Absolutists about torture must also reject these same demands of justice. If one could prevent a man from torturing an innocent child only by torturing the man, absolutists insist that it would be wrong to torture him, even if the torture one would inflict on him would be less bad than that which he would inflict on the child.³ Questions about the just distribution of harm simply do not arise.

I will return to this problem later. Before concluding this section, it is worth noting one further objection to absolutism that is particularly acute in the case of an absolute prohibition of torture. All moral theories have line-drawing problems, but absolutist theories are particularly vulnerable, for they have to draw a line between acts that are absolutely forbidden—impermissible no matter what the alternative might be—and acts that can be permissible. Torture, no matter how it is defined, involves the deliberate infliction of harm. How severe the harm must be to count as torture is of course a question that is much debated, and to which the Bush administration's "Bybee memo" gave a preposterous answer. The important point here, however, is that if the act-type "torture" is supposed to be absolutely impermissible, it must be defined in such a way that it is plausible to say that any act that counts as torture is absolutely impermissible. Absolutism about torture would be intuitively unsustainable if, for example, twisting a person's arm to cause him pain were to count as torture. Indeed, in order for their view to seem at all plausible, absolutists are under pressure to set the threshold for torture rather high. But suppose they are able to define the threshold with some precision, so that the deliberate infliction of any degree of pain or suffering above that threshold counts as torture, provided other relevant conditions are also satisfied. They then face the question: "Why is the deliberate infliction of pain just above the threshold incapable of justification, while the infliction of pain just below it can be permissible, given that the difference between the two degrees of pain is so slight?" I doubt that there is any satisfactory response to this challenge.

2. TORTURE IN PRINCIPLE

In the debate about torture, the notorious "ticking bomb" argument enlists our intuitions against absolutism. This argument deploys the familiar hypothetical example in which we have captured a terrorist who we know has planted a nuclear bomb in a city. The bomb will detonate soon unless we disable it, but the terrorist will not tell us where it is hidden. Our only hope of finding it is to torture him.

If nothing else, this example exposes the intuitive implausibility of absolutism about torture. Opponents of torture are often evasive in addressing the question

whether torture would be morally permissible in this case. I do not, however, think that it aids the credibility of the anti-torture case either to deny that torture would be permissible in this example or to refuse to address the question, as many opponents of torture do. We should concede that torture would be morally permissible, or perhaps even morally required, in this hypothetical case and then ask what implications that concession has for matters of policy and law. I will shortly try to show that advocates and opponents of torture alike tend to exaggerate the significance of the example and to misinterpret its intuitive force.

Opponents of torture tend to argue that the ticking bomb example is unrealistic, as indeed it is. It presupposes a high degree of reliability in the belief that there really is a nuclear bomb that will otherwise detonate, that the person we hold captive planted it, or at least knows where it is, that torture will be effective in getting him to reveal its location, and so on. But pointing out that actual cases have neither the epistemic features nor the all-or-nothing character of the make-believe example leaves it open that actual cases may nevertheless raise similar challenges.

There have been and will continue to be times when people who are attempting to protect innocent people from terrorism will capture a person they reasonably and indeed correctly believe to be guilty of a terrorist atrocity. They will also believe, and not wholly without reason, that by torturing this person they *might* obtain information that they could not otherwise obtain, and that might enhance their ability to prevent other terrorist acts.

These people will want, and need, moral guidance. Could we honestly tell them that they really face no moral dilemma at all, since it should be luminously obvious that to engage in torture would be absolutely impermissible, odious, and barbaric? Would it be illuminating or persuasive to tell them that torture is ruled out because it is disrespectful of human dignity? What if, following our guidance, they were to refrain from torturing their captive, only to discover later that he did indeed have knowledge of an impending terrorist act that subsequently killed thousands of innocent people and that they might have been able to prevent had they tortured him? On what grounds could we reassure them that, even so, it would have been wrong for them to torture him?

There is in fact a good answer to this question but it is not the facile answer offered by absolutism. I will offer this answer at the end of the paper. But before I can state it, I need to say more about the conditions in which torture might in principle be morally justified.

I have claimed that defenders and opponents of torture alike tend to misinterpret the significance of the ticking bomb case. Defenders of torture usually take it to show that torture can be justified as the lesser evil, or that it can have what in law is called a justification of necessity, and opponents of torture often follow them in making this assumption.⁴ The lesser-evil justification is subject to different interpretations. According to the consequentialist interpretation, the intentional infliction of harm is justified whenever it prevents a greater evil, even when

the evil prevented would be only slightly greater than the evil caused. There is, however, a “threshold deontological” interpretation of the lesser-evil justification according to which there are moral constraints against the intentional infliction of harm that can be overridden only when the evil averted is *substantially greater* than the one inflicted. This latter interpretation of the lesser-evil justification is intuitively more plausible than the consequentialist interpretation, but the ticking bomb case is designed to ensure that both interpretations agree that the constraint against torture is overridden by the magnitude of the harm that could be expected to be prevented only by torturing the terrorist.

As I noted, most opponents of torture are not absolutists about the prohibition of killing—even, I suspect, about the prohibition of the intentional killing of the innocent. They accept, in other words, that one or both of these lesser-evil justifications explain certain exceptions to the prohibition of killing. With respect to torture, however, they worry that even the threshold deontological justification affords insufficient protection against torture. For the essential vagueness of the notion of a “substantially greater” evil makes it difficult to challenge the claim by proponents of torture in any particular case that the threshold has been passed—that is, that the evil to be averted is great enough to justify torture. In practice, therefore, the vagueness of this notion tends to vitiate the distinction between the consequentialist and threshold deontological interpretations of the lesser-evil justification. In practice, the lesser-evil justification tends to be almost limitlessly permissive. If the ticking bomb case is understood as supporting the lesser-evil justification for torture, it becomes readily comprehensible why enthusiastic advocates of torture are fond of it, while opponents fear it.

Suppose that in the ticking bomb case the probability of compelling the terrorist to divulge the location of the bomb would be higher if we were to torture his small child before his eyes rather than torture him. A pure lesser-evil justification does not distinguish between torturing the terrorist and torturing his child. Suppose that we could be confident of breaking the terrorist’s will in time either by torturing him or by torturing his child, but that his will would break much sooner if we torture the child. If torturing the child would inflict less suffering overall, despite the fact that this would in effect involve torturing two people rather than one, a pure lesser-evil justification might *require* that we torture the child. That seems to me clearly wrong, though it is testimony to the intuitive force of the threshold deontological version of the lesser-evil justification that if the stakes were high enough in the ticking bomb case, most people agree that it could be *permissible* to torture the child if that offered the best chance of saving the city, which itself, we might suppose, is home to more than a million children who would otherwise be killed.

But of course the stakes have never actually been nearly this high. To the best of my knowledge, there has never been an actual instance of torture that has been justifiable by appeal to a lesser-evil justification with a high threshold for overriding the constraint against the intentional torture of the innocent. Perhaps

there will be such a case in the future. But the mere possibility that such a case will arise is no basis for the formulation of law or policy, both of which have to be focused on the cases that people actually confront. If a ticking bomb case, in which it would be morally permissible to torture a terrorist or his child, were ever actually to occur, people would not look to law or policy or even moral theory for guidance. In these conditions, it would hardly matter what our law or policy might be, and people would not need a moral theory to tell them that torture would be permissible. For people are, as we know, often greatly tempted by torture even in cases in which the stakes are minor in comparison with those in the ticking bomb case. One contingency that we really do not need to worry about is that people will be inhibited by moral scruples from engaging in torture in a ticking bomb case and will thus allow a city to be destroyed.

When I said earlier that people have missed the significance of the ticking bomb case, I meant that they have taken the lesson of the case to be that there can be a lesser-evil justification for torture. While that is true, it is uninteresting, for it is really nothing more than a rejection of moral absolutism. What people have often overlooked is that there is another and better explanation of why it would be permissible to torture the terrorist in that case. This is that the terrorist, by virtue of his responsibility for a threat of wrongful harm to innocent people, has made himself *liable* to be tortured if that is a necessary and proportionate means of preventing his having planted the bomb from killing those people. To say that he is liable to be tortured is to say that torturing him would not wrong him or violate his rights, in the circumstances.

The appeal to liability is a more familiar and less controversial justification for harming people than the appeal to the lesser evil. In criminal law, the infliction of punishment is justified on the ground that the criminal has made himself liable to be punished by virtue of his moral responsibility for a criminal act, usually involving harm to the innocent. In tort law, the imposition of a burden of compensation is usually justified on the ground that the tortfeasor has, through her own fault, made herself liable to compensate the victim or victims of her action. And the best account of permissible defense is that it is justified because the aggressor has made himself liable by virtue of his moral responsibility for a threat of wrongful harm to another. In each case, the justification for the intentional harming of the person who is liable is a matter of justice in the distribution of harm. In criminal law, the usual view is that it is a demand of retributive justice to inflict on wrongdoers the harm that they deserve (even if the *aim* of punishment is to prevent or deter further criminal action). In tort law, it is typically thought to be a matter of corrective justice that harms should be redistributed *ex post* in accordance with people's responsibility for their occurrence. And in the law of self-defense, it is a matter of preventive justice that inevitable harms should be distributed *ex ante* to those who are morally responsible for the fact that others will otherwise be wrongfully harmed.

In the ticking bomb case, the torture of the terrorist could be justified as a matter of preventive justice. Because of his own previous wrongful action and his present wrongful refusal to avert the effects of his earlier act, he is morally responsible for having made it inevitable either that millions of innocent people will be killed or that he will be tortured. Justice requires that what is, for us, an unavoidable harm be distributed to him rather than being allowed to be inflicted by him upon the innocent. While the fact that the harm we inflict is much the lesser of the two evils effectively guarantees that our action is proportionate, it is not a necessary condition of the permissibility of our action. We would be justified in torturing the terrorist even if all we would thereby avert was the equivalent torture of only one innocent person which the terrorist's previous action had made otherwise inevitable. It is, indeed, a commonplace in the theory of justified defense that a person acting culpably can be liable to suffer a *greater* harm than that which the defensive action averts.

Note also that in this latter case involving a choice between tortures, the justification for torturing the terrorist does not extend to the torture of his child. While the terrorist's action has made him liable to be harmed, his child is entirely innocent. The child has done nothing to lose his right not to be tortured as a means of preventing even the more severe torture of another innocent person. Those who reject moral absolutism must concede that the child's right not to be tortured is capable of being overridden, but it is not overridden in this case. Neither is the terrorist's right overridden; rather, the terrorist has forfeited his right not to be tortured as a means of preventing an innocent person from being tortured.

It is also worth emphasizing that the claim here is only that the terrorist is *liable* to be tortured, not that he *deserves* to be. The claim that a person deserves to be harmed in a certain way entails that it is intrinsically good that he should suffer that particular harm. Although I accept that people can deserve to suffer, I do not accept that a person can deserve to be tortured. I do not, however, have a principled account of the upper limits of deserved suffering.⁵

I have canvassed two forms of justification for harming people—that the harming is the lesser evil and that the victim has made himself morally liable to be harmed—and have suggested that most people accept the lesser evil justification in cases in which the harm that is caused is greatly exceeded in magnitude by the harm that is prevented. This extends, in principle, even to the worst forms of torture—for example, most people would accept that it would be permissible to torture one innocent person for a year if this were the only way to prevent a billion innocent people from being tortured in an equivalent way for an equivalent period. The right not to be tortured is thus not absolute because it can in principle be overridden.

One might argue, however, that it is absolute in another sense. It would be absolute in one respect if, even though it could be overridden, it could not be waived, forfeited, or alienated (or some combination of these). Some analyses of what is morally objectionable about torture may suggest that the right not to be

tortured is absolute in the sense that it cannot be forfeited. David Sussman, for example, offers a Kantian explanation of what it is “about torture that sets it apart even from killing, maiming, or imprisoning someone, such that the circumstances that might justify inflicting such harms would not even begin to justify torture.” He argues that the distinctive evil of torture is that it “does not merely insult or damage its victim’s agency, but rather turns such agency against itself, forcing the victim to experience herself as helpless yet complicit in her own violation. This is not just an assault on or violation of the victim’s autonomy, but also a perversion of it, a kind of systematic mockery of the basic moral relations that an individual bears both to others and to herself. . . . The violence of war or police action may injure or insult an agent’s capacities for rational and moral self-governance, but such violence need not make the victim an accomplice in his own violation.”⁶

One who holds this view might accept that while the right not to be tortured can in principle be overridden in extreme circumstances, it cannot be forfeited, so that a person cannot even in principle be morally liable to be tortured. For the claim that a person has forfeited his right not to be tortured, or that he has made himself liable to be tortured, entails that he would not be *wronged* by being tortured. Yet if torture does to a person what Sussman says it does, it may seem that torture must always wrong its victim. If so, the only justification for torture would be the lesser-evil justification. Torture could never be justified on the ground that the victim had made himself liable to be tortured.

This would in a way be a surprising position for a Kantian to adopt. (Sussman himself reserves judgment on whether torture is absolutely impermissible.) For it fails to take people seriously as autonomous and morally responsible agents. If a person has made it the case through his own autonomous choices that the only way to prevent his previous action from killing innocent people is to exploit his vulnerability in order to turn his will against himself, then that may be precisely what his exercise of his autonomous agency has made him liable to have done to him. How can he have a justified complaint if, for example, by refusing to reveal the location of a bomb he has planted, he is freely continuing to make it necessary either to torture him or to allow him to murder innocent people? All he has to do to avoid being tortured, and thus to avoid becoming an accomplice to his own violation, is to do what he is independently morally required to do. If, in these circumstances, he chooses to be tortured rather than to stop himself from killing innocent people, he cannot plausibly claim to be wronged if he is tortured.

As we have just seen, Sussman accepts that torture is harder to justify than killing. This view is nearly universal, yet it is at least *prima facie* puzzling, since even the most intense torture can be less bad for the victim than death, provided that the torture is of some sufficiently limited duration. Death can be worse than great pain and terror, and even worse than experiencing the treachery of one’s own will, when one finds it “expressing the will of . . . a hated and feared enemy”—provided, of course, that the goods of subsequent life would outweigh

these evils inherent in being tortured.⁷ Why, then, do we believe that it is more seriously wrong to torture a person than to kill him?

Sussman offers a Kantian account that is intended, as I understand it, to explain why the violation of respect for a person as an autonomous agent is more egregious in torture than in killing. There are, however, certain contrasts between torturing and killing that may be more accessible to and more readily appreciated by common sense than the rather esoteric account of the wrong involved in torture advanced by Sussman. One of these is that while it is obvious how killing can be defensive, it may seem that torture cannot be. For when a person is in a position in which it is possible to torture him, he must be incapacitated and thus incapable of posing a threat; hence it cannot be possible at that point to defend oneself or others from him. If killing can be defensive while torture cannot, that might explain, at least in part, why we find torture more seriously wrong than killing, even in some cases in which killing would be worse for the victim.⁸

There are several points one might make here. First, it is possible for torture to be literally defensive. Suppose there were a device that could be used from a distance to cause debilitating pain in any person at whom it is directed. Imagine next that one sees a man on the verge of killing an innocent child and that, while one cannot physically restrain him, one can use this device to cause him to suffer pain so severe that he would become incapable of doing anything other than writhing in agony. If one intentionally kept the man crippled with the most intense pain possible for, say, a quarter of an hour, to give the child time to make a complete escape, it would be hard to deny that one was torturing the man as a means of defending the child against him.

The example of defensive torture raises the question whether, if one had two equally effective means of defending the child—killing the man and torturing him with the device—there would be a moral reason to choose killing rather than torture. Intuitively, it seems that the reverse is true: that it would be better morally to use the teletorture device than to kill the man. If this is true, it forces us to recognize a limit to the scope of Sussman's view. The kind of torture I have described, which we can call "purely defensive torture," does not involve hijacking one person's will in the service of another's and thus does not have the distinctively evil feature that Sussman identifies. Neither, for that matter, does punitive torture, or torture inflicted on some as a means of terrorizing and intimidating others. It is possible, therefore, that only interrogational torture, used to elicit information or confession, is objectionable for the kind of reason Sussman identifies, and thus that only interrogational torture can be claimed to be in general more seriously wrong than killing.⁹

That torture can be purely defensive in the way that I have illustrated is of course of limited significance, since few if any actual cases of torture are defensive in this way. Yet it can be argued that even interrogational torture of a fully incapacitated victim can be literally defensive. Suppose a person has initiated a sequence of events that pose a threat to another. If this person is now powerless

to abort this sequence of events, it is plausible to suppose that attacking him or otherwise harming him cannot be literally defensive, even if that would eliminate the threat via some causal sequence that he is incapable of creating by himself. Yet if he retains the ability to stop the threatening sequence but refuses to do so, he can plausibly be regarded as sustaining and therefore continuing to pose the threat he has created.¹⁰ If, for example, one person has administered a slow-acting poison to another and refuses to give the victim the antidote that only he, the poisoner, possesses, it would seem to count as an act of defense if the victim were to kill the poisoner in order to get the antidote from him. If, by contrast, an uninvolved third party was the only person in possession of the antidote and refused to give it to the victim, it would not seem that the victim's killing the third party would count as defensive. Allowing a threatening sequence to continue seems to count as posing the threat only if one has oneself initiated the sequence.

Although these descriptions seem intuitive to me, I assume that many others will say that if the third party does not pose a threat, neither can the poisoner do so now, since both now do the same thing: withhold the antidote. In neither case, therefore, would killing the person to get the antidote be defensive. Suppose this is right. In that case, torturing the terrorist in the ticking bomb case cannot be defensive either.

But whether torturing the terrorist is literally defensive is irrelevant if the moral justification for defense applies equally to the torture of the terrorist. I have argued elsewhere that in most cases of justified defense, the reason that defensive violence is justified is that the person who poses a threat has made himself liable to attack by virtue of his moral responsibility for a threat of wrongful harm to another. Whether he poses the threat now or created the threat through previous action is merely a matter of the timing of the act that makes him responsible for the threat, and that seems irrelevant to his liability.¹¹ (Of course, if the act through which a threat has been created lies in the past, that normally affects the justification for violent action against the agent in various ways. Violent action against an agent now typically can do nothing to avert a threat that was created by his past action. Yet past action is in general a firmer foundation for liability than threatening action in progress. Even though a threat created by past action may never materialize, there may be no uncertainty about whether the act was done, whereas action in progress might be aborted by the agent before any harm is done.)

In spite of all this, there is one difference between interrogational torture and killing in self-defense that may be significant. In killing someone in self-defense, one simply eliminates him as a threat. But interrogational torture involves harmfully *using* the victim as a means of averting a threat, albeit a threat for which he is assumed to be responsible. This difference between merely eliminating a person and opportunistically exploiting him has been identified as presumptively morally significant.¹² One can read Sussman's analysis of the wrong involved in torture as an explanation of the special wrongness of the particular form of using it involves—namely, the use of the victim's will against himself. Yet, as I noted earlier, not all torture takes this

form. What I called purely defensive torture does not involve using the victim at all but in a sense eliminates him, though merely temporarily rather than permanently. So *if* what is particularly objectionable about interrogational torture is connected with the way it uses the victim, we might draw at least two conclusions. First, interrogational torture may be impermissible even when the stakes are high enough that, were it possible to avert the threat by defensive killing instead, the killing would be permissible. This could be true even though the harm to the victim involved in the torture would be significantly *less* than the harm of being killed. Second, purely defensive torture is not objectionable for the special reason that interrogational torture is and therefore might be permissible when the stakes are high enough to justify defensive killing. Indeed, purely defensive torture should be morally *preferable* to defensive killing provided the stakes are high enough to justify killing, and being tortured would be less harmful to the victim than being killed.

I remain skeptical, however, of the suggestion that a person cannot be liable to interrogational torture because it *necessarily* wrongs a person to be used in this way. This does not mean that I think the distinction between eliminative harming and opportunistic using is without moral significance. On the contrary, I am inclined to accept that it is more seriously objectionable to use a person opportunistically than to eliminate a person as an obstacle, if the degree of harm caused and other relevant factors are equal, *and if* the people treated in these ways are relevantly innocent. But the significance of the distinction is vitiated—or at least the distinction cannot by itself make the difference between permissibility and impermissibility—when harm is inflicted on people who are relevantly guilty, or culpable. When the stakes are high enough to justify the defensive killing of a person whose present action culpably threatens the innocent, they are also high enough to justify—in *principle*—the interrogational torture of a person whose *past* action culpably threatens the innocent, at least if the harm caused by the torture would be less than the harm of being killed.

I will conclude this section by noting a point that emerges when we consider the possibility of justifying torture by appeal to the victim's liability rather than by claiming that torture is the lesser evil. Discussions of interrogational torture often focus, quite rightly, on the uncertainties facing those who would practice it, and on the way these uncertainties are blithely stipulated away in hypothetical examples, such as the ticking bomb case.¹³ In actual cases in which interrogational torture might be used to gain information about terrorist activity, the uncertainties and thus the possibilities for mistake are legion. The person tortured might not be a terrorist at all; even if he is, his organization may have no plans for further terrorist activity; even if it does have such plans, he may know nothing about them; even if he knows about them, he may lie, simply saying whatever he judges his captors want to hear, in order to stop the torture; he may die under the stress; and so on.

Of these uncertainties, one is morally more significant than the others. Consider two possible types of case.

(1) We are certain, beyond any possibility of reasonable doubt, that there is a terrorist plot against us, and that an attack is impending. We have captured a person of whom we reasonably believe that there is a significant probability that he is a terrorist and has knowledge that might enable us to prevent the attack. But in fact this person is not a terrorist and has no relevant knowledge.

(2) We have captured a person who we are certain is a terrorist. (Suppose that there are videos, taken independently by unrelated observers, of this person throwing a grenade into a school bus filled with children, and that we later subdued and captured him as he was entering a crowded restaurant with bombs strapped beneath his overcoat.) We reasonably believe that there is a high probability of an impending terrorist attack by members of his group and that he has knowledge of the plot. But in fact (2i) there is no plot, or (2ii) while there is a plot, he has no knowledge of it.

Suppose that, in both cases, we torture the captive in an unavailing effort to gain information. In both cases, our action is objectively wrong, for we have tortured a person without any possibility that something good could come of it, though we could not have known this in advance. In both cases, it is possible that our action is *subjectively* permissible, in the sense that if our factual beliefs, which I have stipulated are reasonable or epistemically justified, were all true, then our action would be objectively permissible. Our action might be subjectively permissible if in the first case the probability of an impending, large-scale attack were very high, or if in the second case the probability that our captive is a terrorist with knowledge of the impending attack were very high.

There is nevertheless an important difference between the cases that makes it significantly more difficult to justify interrogational torture in the first case than in the second. This is that in the first case our action clearly wrongs the victim, or infringes his rights, whereas that may not be true in the second. In the second case, our captive has freely acted in ways that have now created a situation in which we reasonably believe that we must choose between torturing him and allowing a large number of innocent people to remain at significant risk of being killed by action in which he is complicit and for which he would therefore be jointly responsible. In reality, our epistemically justified belief that we face this dilemma is false. But it is the terrorist's fault, not ours, that we are in this situation. By his own culpable action, he is responsible for our justified, though false, belief that he continues to pose a threat to innocent people. He cannot reasonably expect us to accept his assertion that he has no knowledge of any further plot. He has therefore imposed on us the subjective necessity of acting in the absence of relevant knowledge. In these conditions, he has no justified complaint if we choose to try to reduce what we reasonably perceive to be the great risks that he and his confederates pose to numerous innocent people by inflicting grave harm on him.

3. TORTURE IN PRACTICE

Thus far I have argued that interrogational torture can in principle be morally justified in a way that is continuous with the primary justification for self-defense and defense of others. But having made this concession, I will now argue that it is of virtually no practical significance.¹⁴ Whether torture can be morally permissible is less significant as a question of individual or personal morality than it is as a question of institutional morality—that is, the moral principles governing the design and functioning of social institutions. This is not to deny that the question whether it is morally permissible to participate or engage in torture arises with considerable urgency for some individuals. But I suspect that the vast majority of those who are in a position in which this question might arise are not much interested in morality and are thus disinclined to consider the question at all. For the minority who may wrestle with the question, deliberation is likely to be conducted principally by reference to the law; that is, they will look to the law for moral guidance. And in any case, the fact that interrogational torture is not a private activity but a political one means that morality must govern the practice not primarily through appeals to individual conscience but by dictating what law and policy should say about it.

What, then, does morality imply about how the law should treat the practice of torture? In conditions in which we could expect full compliance with the law of torture but not with other areas of the law, or with morality, the law should of course permit torture on those rare occasions when it would be morally justified—that is, when the victim is liable, the stakes high, and the uncertainties minimal—and prohibit it in all other cases. But these are obviously not the conditions in which we live. In the conditions in which it is our misfortune to live, a law that would simply restate the permissions and prohibitions of morality would be wholly infeasible. In these conditions, state officials contemplating the use of torture are their own judges, and those whose goals are unjust are likely to believe that they are just. And even when they are aware that their goals are unjust, they are unlikely to have scruples about means and will claim moral justification whenever torture seems expedient. Even those whose goals are just will be tempted to perceive or to concoct a moral justification when none exists.

If we could give a precise account of the conditions of moral justification for interrogational torture and could effectively enforce a law that simply prohibited torture in all cases in which those conditions were not met, so that all those who used torture in the absence of moral justification could expect to receive punishment, then such a law might be practicable. But even if we could produce a determinate set of conditions in which interrogational torture would be morally justified, a law that permitted torture only in those conditions would not be enforceable. States would shield their own torturers and states themselves, or at least the more powerful ones, would be shielded by our general inability to bring effective sanctions against them.

It seems, therefore, that if we grant any legal permission to use torture, particu-

larly one that attempts to capture the complex conditions of moral justification, it will be exploited by those whose aims are unjust and either abused or interpreted overly generously even by those whose aims are just. Throughout human history, torture has been very extensively employed, but the proportion of cases in which the use appears to have been morally justified seems almost negligible. Part of the reason for this is that morally decent people are naturally repelled by the practice of torture and are reluctant to use it; thus it tends to be used far more frequently by those who are both unjust and cruel.¹⁵ This does not mean that it is uncommon among peoples that subject themselves to democratic constraints. What has been called “clean torture”—torture that leaves no marks—has been employed by Western, democratic states far more often than most of us suspect.¹⁶ But this brings out another important point, which is that the forms of torture used by undemocratic states tend to be even more hideous than the “clean” forms favored by states with provisions for democratic accountability. The tortures inflicted at Abu Ghraib and Guantanamo are in general (at least so far as we know at present) quite tame compared to the techniques used, for example, by the fascist regimes in Latin America that the U.S. supported during the 1970s and 1980s—though these regimes were, admittedly, more interested in torture for terrorist rather than interrogational purposes, and so were free to be as imaginative as they liked.

The crucial points are these. When torture has been practiced, it has been unjustified far more often than it has been morally justified. In part this is because it is more often used by the unjust against the just than by the just against the unjust. The forms that it takes in the hands of those whose aims are unjust tend, moreover, to be the most horrible forms imaginable. It therefore seems that anything that makes it easier for governments to use torture is almost certain to have terrible effects quite generally, and in particular to result in far more violations of human rights than would otherwise occur. Any legal permission to use torture, however restricted, would make it easier for governments to use torture, and would therefore have terrible effects overall, including more extensive violations of fundamental human rights. The legal prohibition of torture must therefore be absolute.

This may strike most of us as plausible in the case of international law. Few of us, after all, would like to see loopholes that could be exploited by regimes such as the former Ba’athist government in Iraq. But some people, known as “exceptionalists,” argue that the U.S. is different and that we can safely have highly circumscribed provisions for the legal use of torture without precipitating the widespread practice of torture by vicious and undemocratic regimes, which will probably use it to the extent that they find it expedient no matter what we do. But this is a delusion. The Bush administration has provided ample proof, if any were needed, that we cannot be trusted to use torture only on those very rare occasions on which it would be morally justified. More importantly, we cannot proceed with torture the way we have with nuclear weapons—that is, by permitting it to ourselves while denying it to others by means of security guarantees, economic rewards, and other measures

designed to make abstention in the interests of all. If we permit ourselves to use torture, we thereby forfeit any ability we might otherwise have to prevent its use by others. Any efforts we might make would be no more effective than a proselytizing defense of vegetarianism by someone complacently enjoying a steak. Our only hope of being able to impose legal and other constraints on the use of torture in the service of unjust ends by vicious and cruel regimes is to deny the option to ourselves as well, even in cases in which we believe it would be permissible.

If we are to deny ourselves the option of torture, we must reject it not only legally but institutionally. We must make it transparent to external observers that we do not train our interrogators in techniques of torture, do not permit them the use of special equipment for torture, and will hold them liable to harsh punishments if they ever do use torture, even with higher authorization.

A total legal prohibition of torture, both domestically and internationally, will not, of course, prevent its use. But it can make it costlier for governments to practice torture, and anything that makes torture harder to practice is important.

It is also obvious that a legal prohibition of torture does not preclude an effective defense against terrorism. I think we should concede that there may be occasions on which obedience to a law prohibiting interrogational torture will make innocent people more vulnerable than they would otherwise be. But that is compatible with its being the case that we and others will be more secure overall if, in an effort to eliminate torture altogether, we refrain from using it even when it would in fact help us, at least in the short term.

There is an analogy here with an effective policy of gun control. If we could be largely successful in eliminating the private possession of handguns, we would, in general, be substantially more secure than we are with widespread private possession, even when most people's motives for keeping a gun are defensive. It is true that effective gun control would leave some guns in the hands of criminals and that there would be occasions when the policy would deny the most effective means of self-defense to a person confronted by an armed criminal. But we should simply accept the inevitability of those occasions as the price of a policy that would greatly reduce the occasions when self-defense would be necessary, thereby greatly enhancing people's security overall. It would be irrational to prefer a more effective means of defense in the event of an attack, if the cost were that one would be more likely to be attacked, and therefore far more likely to be killed than if one were denied the more effective defense.

At the end of the first section of this paper I raised the question what we could say to people who have refrained from torture only to find that if they had used it they could probably have averted a tragedy. On most occasions—that is, in cases in which the certainties about threat, liability, effectiveness, and so on that characterize the ticking bomb case are absent—what we can say to them is this:

What you did was *subjectively* right—that is, it was what you ought to have done given the beliefs you reasonably held at the time. There was no rational

basis available to you for doing otherwise than you did. In the great majority of situations epistemically indistinguishable from the one you were in, torture would have been unnecessary or ineffective and thus would have been objectively wrong. You were simply unlucky that your reasonable beliefs turned out, improbably, to be mistaken. If you could have known all the facts, it *might* have been permissible for you to use torture despite the effect that might have had in eroding respect for the taboo against torture that we must continue to work to establish. But the level of certainty about the relevant facts that would have provided that justification was simply not available to you. In the circumstances in which you had to act, you did exactly as you ought to have done.

There are many objections to the argument of this paper that I cannot address in the space allotted to me here. But I will conclude by noting and briefly responding to one. One might accept that moral absolutism about torture is mistaken and yet believe that people generally will be more likely to repudiate the use of torture if they believe that it is absolutely prohibited by morality than if they believe that it can sometimes be permissible. If that were true, it is arguable that morality itself would require that we try to deceive ourselves and others into accepting the absolutist position. I am reasonably confident that a world without torture but in which people held mistaken absolutist beliefs would be better than a world in which people held the view for which I have argued but were insufficiently motivated by it, so that torture continued to be used. But I do not think that we face this choice. I think the case I have advanced against torture is in fact quite strong. It is simple without being simple-minded, and its simplicity makes it accessible and frees it from reliance on rhetoric for its impact. Indeed, I think it will actually be more convincing than the absolutist position to ordinary people, whose modes of thought tend to be more receptive to pragmatic considerations than to high-minded moral doctrines that they may find more suited to guiding the conduct of saints than to determining the policies of states.¹⁷

Rutgers University

NOTES

1. Ben Juratowitch, "Torture is Always Wrong," *Public Affairs Quarterly*, vol. 22, no. 2 (April 2008), pp. 81–90.

2. There are many examples, of which Juratowitch is only one. Another is Kim Lane Scheppele, who writes, "I do in fact believe that torture is always and absolutely wrong, given the position we should accord to human dignity, even that of terrorists." She does not, however, argue for this view but instead develops a strong pragmatic case against the practice of torture. See Kim Lane Scheppele, "Hypothetical Torture in the 'War on Terrorism,'" *Journal of National Security Law and Policy*, vol. 1 (2005), pp. 285–340, p. 287.

3. See Jeff McMahan, "Torture, Morality, and Law," *Case Western Reserve Journal of International Law*, vol. 37 (2006), pp. 241–248, pp. 243–244.

4. For example, Juratowitch (p. xxx) takes the ticking bomb case to support "the consequentialist case for torture," and Scheppele (p. 293) contends that "hiding behind this hypothetical is an implicit consequentialist argument that torture would be justified if the consequences of *not* torturing were serious enough."

5. I am waiting for Shelly Kagan's much-anticipated book, *The Geometry of Desert*, to discipline and instruct my unruly intuitions about desert.

6. David Sussman, "What's Wrong with Torture?" *Philosophy and Public Affairs*, vol. 33 (2005), pp. 1–33, pp. 3 and 30.

7. *Ibid.*, p. 29.

8. In his seminal paper, "Torture" (*Philosophy and Public Affairs*, vol. 7 [1978], pp. 124–143), Henry Shue argues that one of the main reasons why torture is repugnant is that it involves the use of violence against someone who is entirely helpless or defenseless. One way of interpreting the significance of this claim is that the use of violence against a person who is helpless cannot be defensive.

9. In his final section (pp. 30–33), Sussman seeks to extend his analysis of the wrongness of torture so that punitive torture turns out to be distinctively objectionable for reasons similar to those that explain the wrongness of interrogational torture. I find the claims of this section less plausible, and am skeptical that these claims could show, for example, that punitive torture is more seriously objectionable than capital punishment.

10. This point is suggested in Jeff McMahan, "The Ethics of Killing in War," *Ethics*, vol. 114 (2004), pp. 693–733, p. 721, note 46. Sussman independently makes the same point on p. 16.

11. Jeff McMahan, "The Ethics of Killing in War"; and Jeff McMahan, "The Basis of Moral Liability to Defensive Killing," *Philosophical Issues*, vol. 15 (2005), pp. 386–405.

12. Warren S. Quinn, "Actions, Intentions, and Consequences: The Doctrine of Double Effect," *Philosophy and Public Affairs*, vol. 18 (1989), pp. 334–351, p. 344.

13. See, for example, Scheppele, "Hypothetical Torture in the 'War on Terrorism'"; David Luban, "Liberalism, Torture, and the Ticking Bomb," *Virginia Law Review*, vol. 91 (2005), pp. 1425–1461; and Henry Shue, "Torture in Dreamland: Disposing of the Ticking Bomb," *Case Western Reserve Journal of International Law*, vol. 37 (2006), pp. 231–239.

14. Here I am in agreement with Luban, Scheppele, Shue, and many others.

15. In the immediate aftermath of World War II, General Douglas MacArthur negotiated a secret amnesty for the Japanese perpetrators of atrocities involving medical experimentation on captives. He offered them immunity from prosecution for war crimes in exchange for exclusive access to the data derived from the experiments, which he judged would be useful to the U.S. military in future conflicts. What is noteworthy here is that he judged even so shameful an act as this to be preferable to obtaining the Japanese files by means of torture. Thus, in acquiescing in the exchange, he remarked to a subordinate: "Well, if you feel that you cannot draw out the information, we are not given to torture." See Hal Gold, *Unit 731 Testimony* (Tokyo: Tuttle Publishing, 1996), p. 97.

16. See Darius Rejali, *Torture and Democracy* (Princeton, N.J.: Princeton University Press, 2007).

17. I am immensely grateful to Jamie Mayerfeld for pressing me relentlessly to be attentive not only to the validity of my arguments but also to the ways in which they might be misinterpreted or misused. Although he will still disagree with it, this paper is much better for my exposure to his passionate concern about this issue.